



HAL
open science

Harnessing the geographical flexibility of distributed computing clouds for cooperative self-consumption

Benjamin Camus, Anne Blavette, Fanny Dufossé, Anne-Cécile Orgerie

► **To cite this version:**

Benjamin Camus, Anne Blavette, Fanny Dufossé, Anne-Cécile Orgerie. Harnessing the geographical flexibility of distributed computing clouds for cooperative self-consumption. ISGT-Europe 2018 - 8th IEEE PES Innovative Smart Grid Technologies Conference Europe, Oct 2018, Sarajevo, Bosnia and Herzegovina. pp.1-6, 10.1109/ISGTEurope.2018.8571867 . hal-01856665

HAL Id: hal-01856665

<https://hal.science/hal-01856665>

Submitted on 13 Aug 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Harnessing the geographical flexibility of distributed computing clouds for cooperative self-consumption

Benjamin Camus^{*}, Anne Blavette[†], Fanny Dufossé[‡] and Anne-Cécile Orgerie^{*}

^{*} Univ Rennes, Inria, CNRS, IRISA – Rennes, France – Email: {benjamin.camus, anne-cecile.orgerie}@irisa.fr

[†] Univ Rennes, SATIE, CNRS, Rennes, France – Email: anne.blavette@ens-rennes.fr

[‡] LIG, Inria, Grenoble, France – Email: fanny.dufosse@inria.fr

Abstract—Smart grids allow to efficiently perform demand-side management in electrical grids in order to increase the integration of fluctuating and/or intermittent renewable energy sources in the energy mix. In this paper, we consider a distributed computing cloud partially powered by photovoltaic panels as a self-consumer that can also benefit from geographical flexibility: the computing load can be moved from one data center to another one benefiting from better solar irradiance conditions. The various data centers composing the cloud can then cooperate to better synchronise their consumption with their photovoltaic production. The simulation-based results show that our approach can significantly increase the self-consumption of renewable energy and improve the state-of-the-art management techniques for energy-efficient distributed clouds.

Index Terms—cooperative self-consumption, spatial flexibility, computing clouds

I. INTRODUCTION

Demand-side management is key for increasing the level of renewables in the energy mix. This concept implies modifying the consumer’s energy consumption by means of different incentives (e.g. economic) with respect to the fluctuations of the renewable electricity production in order to maintain the safe and stable operation of the power system. So far, a number of flexible load types has been extensively studied for demand-side management applications, such as controllable water heaters, air-conditioning equipment, refrigerators, etc. as well as electric vehicles. Traditionally, flexible loads are characterized by a certain degree of temporal flexibility. This means that their energy consumption can be reduced partially or entirely (i.e. interruptible loads) during a certain amount of time, usually with a rebound effect, or that their activation can be postponed (i.e. deferrable loads).

However, it may also be interesting to consider spatial flexibility. Spatial flexibility can be defined as the ability of a load to migrate physically from one node in the electrical network to another in a sufficiently short amount of time to be relevant for demand-side management. Spatial flexibility may be used for addressing local network issues, such as line congestion or voltage control. It is important to note that one approach to solve the mentioned local network issues consists in curtailing the renewable electricity generation, thus losing energy. Hence, spatial flexibility can represent an alternative to curtailment. Another advantage of spatial flexibility consists in accompanying the spatial fluctuations inherent to variable renewable electricity generation from wind

or photovoltaic sources. In other words, such loads could be expected to migrate from one region to another of the power system if the electricity generation from renewables becomes more important in the latter than in the former. This represents a relevant manner to further harness renewables, in complementarity with temporal flexibility.

Spatial flexibility can be provided by loads such as distributed computing clouds. Clouds include a large number of computing servers (gathered in data centers) which can be virtualized, thus offering on-demand virtual machines (VMs) to end-users that are adapted to their needs in terms of computing power. These VMs can be deployed in any data centers, and can be displaced from one to another, which is usually referred to as “VM migration”. In particular, these VMs can migrate from a data center where renewable electricity generation is becoming scarce to another with more favorable conditions.

Hence, spatial flexibility can represent an interesting approach for cloud managers, as they are becoming more and more aware of the increasing financial and environmental impact of their data centers ever growing energy consumption [1]. Hence, more and more data centers are being equipped with photovoltaic (PV) panels in the perspective of self-consumption. VM migration can thus allow to increase the cloud self-consumption, while not impacting the quality of service provided to the cloud customers.

Although a large number of studies is available on VM migration for better harnessing renewables [2], [3], none has compared this approach with the alternative which consists in exchanging photovoltaic energy between the data centers through the electrical network.

Considering that data centers are supposed to be relatively distant from one another, in order to harness renewables spatial diversity, the energy exchange between them is virtual rather than physical. In other words, it is considered that the data centers can exchange energy as part of a collective self-consumption scheme.

Collective self-consumption policies and decrees have been recently adopted or are under development in particular in Europe [4], and especially in countries such as France [5] and Germany [6]. This concept consists in allowing multiple consumers and renewable energy producers, all located in a small geographical area, to establish power purchase agreements (or similar schemes) between themselves. Legislation differs between countries, in particular regarding the definition

of the geographical area which can span from a single site (e.g. a building equipped with PV panels on its roof) to a wider, but still local, geographical area where consumers and producers are connected to the same low-voltage feeder. The collective self-consumption scheme proposed by the authors in this paper is inspired from existing or under development policies, but it represents an extended, multi-site version. This version is deemed to require the establishment of only relatively simple measures, facing no deep technical challenges, and able to unlock the potential of spatial flexibility inherent to distributed computing clouds.

However, it is important to take into account that absorbing energy from, or injecting energy into, the grid leads to Joule losses, transmission and distribution equipment aging, etc. as well as grid operator management costs. This impact can be taken into account as a cost based on the network use tariff. This tariff is usually calculated based on the amount of energy and power transmitted by, or to, a grid-connected installation through the grid. Hence, self-consuming the PV electricity generated locally, i.e. not absorbed from, or injected into, the grid, reduces this cost. On the contrary, energy exchange between data centers, or between data centers and other customers or power plants increases this cost. On the other hand, migrating a VM requires an additional amount of energy (compared to exchanging energy through the electrical network), which also has a cost, as it will be detailed in Section II. Hence, this paper will focus on the comparison in terms of costs between the VM migration and the energy exchange approaches, and a combination of these.

Section II provides some background on Cloud computing. The hypotheses, models and the self-consumption optimization algorithm are detailed in Section III. Finally, the results will be described in Section IV, and future work is presented in Section V.

II. BACKGROUND ON CLOUD COMPUTING

Distributed cloud infrastructures consist of several data centers linked through dedicated telecommunication networks. Each data center comprises servers hosting virtual machines that run end-user applications, such as online services (e.g. emailing, file sharing). Upon creation, a virtual machine (VM) is allocated with virtual resources: memory (RAM), computing power (CPU) and disks. These allocated virtual resources are mapped to the physical resources of the server hosting the VM. Depending on its physical characteristics, a given server can host several VMs at the same time. Using its computing resources increases the power consumption of the server. Yet, servers also consume power while being idle but power on. This idle power consumption can reach 50% of the maximal power consumption of recent servers, thus justifying energy saving policies based on server shutdown techniques [7]. These policies consist in consolidating VMs, i.e. in gathering them on a fewer number of servers, so that a reduced number of the latter remains on, while the unused servers are switched off. Displacing virtually one VM from one hosting server to another is usually referred to as “migration”, as mentioned

earlier [8]. However, these operations – switching on/off servers and migrating VMs – consume time and energy. Cloud managers should thus carefully optimize their use of these operations as well as the initial allocation of VM onto servers. As the electricity bill of data centers keeps increasing with the success of Internet-based services, Cloud providers resort more and more to on-site renewable energy sources [9], [10].

III. HYPOTHESES, MODELS AND METHODOLOGY

In this section, we introduce the considered hypotheses and models for the flexible consumer (distributed cloud) and the smart grid. Then, we describe our methodology to optimize self-consumption in a cooperative manner between the cloud data centers.

A. Consumer model: distributed cloud

1) *Cloud workload*: Clouds are supplying computing power by the means of VMs allocated upon request to end-users at any time. When receiving a user request, the cloud manager decides to which data center (DC) and to which server the VM will be allocated.

We use the Eucalyptus Cloud traces of [11]. This workload combines the traces of six different real cloud systems. It includes the list of VM arrival times, required numbers of cores (i.e. virtual CPU allocated to VMs) and execution lengths. This workload is not a favorable scenario with respect to photovoltaic self-consumption, as it does not follow a day/night pattern, that is with a reduced workload during night time when photovoltaic electricity generation is null. Such an unfavorable workload is typical of Infrastructure-as-a-Service Clouds [11], thus representing a realistic scenario.

2) *Telecommunication network*: Once a VM is allocated to a server, it can be re-allocated to another server (in the same DC or not) during its lifetime without service interruption by using VM live-migration mechanisms [8]. Such a migration requires using the telecommunication network linking the two concerned servers (host and destination). The bandwidth capacities of telecommunication links limit the number of VM migrations that can be executed simultaneously, thus limiting the spatial flexibility of the cloud load.

In order to use a realistic distributed cloud infrastructure, in our experiments, we use the characteristics of the French experimental testbed Grid’5000 [12]. Thus, the studied cloud comprises a total of 1,035 servers spread across 9 DCs in Luxembourg and 8 different cities in France (Grenoble, Lille, Lyon, Nancy, Nantes, Rennes, Sophia and Toulouse). The DCs are linked together thanks to 10 GBps Ethernet links. Inside each DC, the servers use 1GBps Ethernet links.

3) *Server power model*: The characteristics of each server are based on the Taurus servers of Grid’5000. Each server is equipped with 2 Intel Xeon E5-2630 CPU with 12 cores in total, 32GB memory and 598GB storage. In order to determine the power consumption of each node, we implement the power model of [13] that is based on real measurements made on Taurus nodes. These measurements notably state that a Taurus server consumes 8W when powered off, 97W

when idle, and 220W at 100% CPU load. Moreover, this model provides also the energy consumption and duration corresponding to powering off and powering on a server (average values over 10 real measurements). These values are used in our simulations as well as the relation between CPU load and power consumption for non-idle servers. Similarly, we use the migration model provided in [14] to determine for each VM how much time it takes to be migrated, and consequently the energy consumption on both servers (the one initially hosting the VM and the one receiving the VM after migration).

4) *On-site electricity generation*: Each DC produces its own renewable electricity thanks to photovoltaic (PV) panels. It is important to note that the PV electricity generation is considered as stochastic, as it can only be predicted with a certain margin of unknown error. Also, the PV electricity generation is highly dependent on the meteorological conditions at each DC. The PV production forecast is modeled as a normal law truncated in 0 in order to compute the expected production and to optimize VMs allocation. When the local renewable electricity generation is insufficient with respect to the DC consumption, the additional energy provider is in charge of powering the DC.

We use real recordings of PV panels power profiles collected as part of the Photovolta project [15] carried out at the University of Nantes, France. The data considered here corresponds to the power produced by 4 Sanyo HIP-240-HDE4 PV panels at a time resolution of five minutes over one week. In order to have heterogeneous trajectories between DCs, we select recordings starting at different dates. Thus, using these real traces we take account of the temporal and spatial variations of the solar irradiance, and therefore of the intermittent nature of PV panels electrical production. For each DC, the number of PV panels is dimensioned to have one PV panel for 3 servers and PV signals are scaled accordingly.

5) *Cloud manager*: The centralized cloud manager deals with incoming VM requests by end-users and controls the servers in each DC. It is responsible for VM allocation and migration decisions, and for switching on/off servers. While VMs can be submitted by users at any time, for optimization purpose, our allocation algorithm (detailed in Section III-C) is executed periodically. We choose a time step of 5 minutes as photovoltaic traces are provided with this granularity. It means that every 5 minutes, the cloud manager launches the VM allocation process and deploys the new VMs requested during the previous 5 minutes.

B. Smart grid model

The cloud acts as an aggregator of self-consuming data centers which are supplied either by themselves or by another energy supplier, if necessary, as allowed by collective self-consumption rules.

A data center is considered to consume its local PV electricity in priority. Then, it injects the excess of PV electricity into the electrical network, if any. Otherwise, it consumes electricity from the electrical network. This electricity will

be considered as exchanged from another data center if one or more data centers inject PV electricity during the same time slot. Otherwise, it will be considered as supplied by the additional energy supplier. The sum of the PV electricity injections into the grid and exchanged with other data centers is considered to constitute a “virtual energy pool” into which data centers in deficit of such electricity can tap. In the absence of energy storage equipment, the remaining amount of PV electricity, if any, is deemed to be sold and consumed by other electricity customers during this time interval. Time slots Δt_g of 10 minutes will be considered. This duration corresponds to the smallest temporal resolution of load curves which is provided to customers connected at the distribution level by smart meters in France [16]. The amount of energy exchanged between data centers through the virtual pool is thus defined as follows:

$$E_{pool}(t_g) = \min \left\{ \begin{array}{l} \sum_i (\max(0, E_{PV_i}(t_g) - E_{DC_i}(t_g))) \\ \sum_i (\max(0, E_{DC_i}(t_g) - E_{PV_i}(t_g))) \end{array} \right\}$$

where $E_{PV_i}(t_g)$ is the PV generation on DC_i at time slot t_g , and $E_{DC_i}(t_g)$ the local energy consumption. The local self-consumption ratio $R_{SC}^L(t_g)$ over each time slot Δt_g is calculated as the ratio of the PV energy consumed locally by DCs by the total amount of energy consumed by the cloud, such as:

$$R_{SC}^L = \frac{\sum_{t_g} \sum_i \min(E_{PV_i}(t_g), E_{DC_i}(t_g))}{\sum_{t_g} \sum_i (E_{DC_i}(t_g))}$$

Similarly, the cooperative self-consumption ratio is the ratio of PV energy consumed by the cloud, including the virtual pool, by the total amount of consumed energy:

$$R_{SC}^C = \frac{\sum_{t_g} (E_{pool}(t_g) + \sum_i \min(E_{PV_i}(t_g), E_{DC_i}(t_g)))}{\sum_{t_g} \sum_i (E_{DC_i}(t_g))}$$

1) *Energy cost*: The energy cost C_E is composed of a fixed and of a variable share. The fixed share $C_{E,f}$ consists in standing charges which have been estimated to 45.79€/month, based on data available in the case of a small industrial customer powered by French electricity supplier EDF [17]. The variable share $C_{E,v}$ of the energy cost (in €/kWh) was estimated based on available data on the kWh unit price paid by residential customers in France, in the absence of data specific to data centers. A fixed electricity price has been considered for the electricity purchased from the additional energy supplier, as proposed by many energy suppliers [18], [19], contrary to a variable price indexed on the energy market prices. The kWh unit price considered in this paper is equal to 0.15€/kWh.

However, a data center may also present an excess of PV electricity which it injects into the grid. This energy can either be exchanged with other data centers through the virtual pool or not. In the first case, the energy exchange is considered to be realised at no cost for the cloud. In the latter case, the PV electricity excess is considered to be sold at a tariff of 0.06 €/kWh, as described in French decree of May, 9 2017 for PV power plants whose peak rated power is less

than 100 kWp [5], as it is the case here. These rules can be expressed as: $C_E = C_{E,f} + C_{E,v}$ where the variable share $C_{E,v}$, corresponding to the amount of energy ΔE , can be expressed as:

$$C_{E,v} = \begin{cases} -0.06\Delta E & \text{if } \Delta E \text{ is injected into the grid} \\ 0.15\Delta E & \text{if } \Delta E \text{ is purchased} \\ 0 & \text{if } \Delta E \text{ is exchanged between DCs} \end{cases}$$

The cost C_E is calculated over each time slot of 10 minutes.

2) *Network use tariff*: As mentioned earlier, power transmission in the electrical network leads to Joule losses, equipment aging, etc. as well as grid operator management costs. This impact can be taken into account as a cost represented by the network use tariff. This cost is paid to grid operators (TSOs and DSOs) by energy suppliers. However, it is eventually paid by energy customers as their bills include a contribution to this charge. Hence, in this paper, the network use tariff was considered to be included in the energy cost when this energy is not exchanged between data centers. However, it is necessary to define this network use cost when it is the case.

This cost was estimated based on the French network use tariff called TURPE (standing for ‘‘Public Electrical Network Use Tariff’’) [20]. This tariff follows four main rules: it is independent of the location of the grid-connected installations and of the distance between the electricity injection and absorption points. However, it depends on the contract power and on the amount of energy consumed. It is important to re-emphasize that the amount of energy on which the network use cost is calculated is restricted to the energy exchanged between data centers through the virtual pool, as the energy cost includes implicitly the network use cost in other cases. Finally, the network use tariff depends on the period j considered, which is usually defined based on the season, the day in the week, and the hour in the day.

The energy E_j that is consumed by the cloud (and originating from the virtual pool) over period j can be expressed as: $E_j = \sum_{t_g \in j} E_{pool}(t_g)$

The network use tariff is composed of a fixed share and of a variable share, in addition to penalties if the contract power is exceeded. The fixed share corresponds to costs related the network management and was estimated at 397.92€/year. Considering that the actual maximal power absorbed from, or injected into, the grid never exceeds the contract power level, no penalties are applied in this case. As regards the variable share C_{N_v} , it depends on both the contract power for all periods and on the amount of energy consumed, and it can be expressed as: $C_{N_v} = \sum_i bS_i + \sum_j^4 c_j E_j$ where S_i is the contract power of data center i , b is the power coefficient, and c_j is the energy coefficient over period j . Four periods j are considered: peak hours and season (period 1, as indicated in Table I), off-peak hours and season peak (period 2), peak hours and season valley (period 3), off-peak hours and season valley (period 4). In addition, two tariff options are proposed: ‘‘long use time’’ and ‘‘short use time’’. As data centers operate 24/7, the ‘‘long use time’’ option was selected. The corresponding numerical values chosen for coefficient c_j

TABLE I
NUMERICAL VALUES OF THE ENERGY COEFFICIENT c_j .

	Period 1	Period 2	Period 3	Period 4
c_j (c€/kWh)	4.18	2.81	1.89	1.74

can be seen in Table I. Coefficient b was calculated as equal to 18.34€/kVA/year.

C. Optimization method

Our approach is named SCORPIUS: Self-Consumption Optimization of Renewable energy Production In distribUTed cloudS. It relies on the previously described virtual pool and costs, and is intended to increase cooperative self-consumption between the data centers. It also employs a probabilistic model for the photovoltaic electricity generation, as mentioned in Section III-A4, that allows to approximate the expected PV production on each DC at any time slot. SCORPIUS can be divided in four main steps.

In the first step, it allocates incoming requested VMs to servers on DCs according mostly to the availability of servers. This algorithm mainly focuses on allocating the VMs on the least possible number of servers. It relies on a Best-Fit heuristic which is a well known algorithm in the computing cloud community [21]. Second, SCORPIUS revisits VM allocation decisions by considering possible reallocations to different DCs based on the predicted PV electricity resources at each data center. More precisely, the VMs pre-allocations are modified by using a Best-Fit heuristic based the expected PV production at each DCs.

Third, SCORPIUS compares the cost of migrating VMs from DCs with insufficient photovoltaic electricity generation to other DCs with more favorable solar conditions, with the costs corresponding to exchanging the considered amount of energy between data centers through the virtual pool, or to purchasing it from the additional energy supplier. It is important to note that migrating VMs requires an additional amount of energy compared to the case when energy is transmitted through the electrical network. This is explained by the fact that, during the migration phase of duration T_m , a given VM runs on two DCs (the hosting and the receiving DCs). Hence, electrical power is consumed by both DCs during this period.

Due to bandwidth constraints, a DC can only migrate VMs one by one. Thus, the amount of VMs a DC can send in a single time slot is bounded by the sum of migration times. In addition, both DCs have to be synchronized during the migration time. Hence, if many DCs send VMs to the same DC, they can not communicate simultaneously.

Figure 1 illustrates this communication issue and shows how our algorithm works with 3 DCs (DC1, DC2 and DC3) sending 4 VMs to 2 other DCs (DC4 and DC5) during the same time slot. First, all DCs are sorted by their expected remaining photovoltaic energy for the rest of the day. The idea here consists in sending VMs from the DCs with the lowest expected remaining photovoltaic energy to those with the highest one. In Figure 1, DC1 starts sending one VM to




Fig. 1. Example of migration plan with 5 DCs.

DC5. When the migration is completed, DC1 has no longer any VM running on it, so it can be switched off. In parallel to this, a VM of DC2 is migrated to DC4. When the migration from DC2 to DC4 is completed, the migration of the VM from DC3 to DC5 can start. In parallel to this migration, when the migration from DC1 is terminated, the second VM from DC2 can be migrated to DC4. This migration plan, resulting from our algorithm, is presented in Figure 1.

Finally, SCORPIUS executes individually on each DC a consolidation algorithm using intra-DC migrations and based on a dichotomy on the number of servers in order to turn off as many servers as possible.

IV. RESULTS

A. Test process

Our approach is tested based on numerical simulations using the SimGrid framework which is dedicated to the simulation of distributed IT systems [22]. SimGrid embeds sound simulation models of CPUs, TCP/IP networks, VMs, and servers energy consumption, which have been assessed both theoretically and experimentally. Thanks to this platform, we accurately simulate the resources usage (i.e. CPU and bandwidth sharing), the execution time and the cloud energy consumption.

We compare SCORPIUS performance against the round-robin algorithm that is traditionally implemented in cloud schedulers. This latter algorithm distributes the workload uniformly among the DCs by allocating incoming VMs to the DCs following a circular order. However, contrary to our approach, round-robin does not take into account the switching off unused servers, nor VMs migration.

B. Results analysis

1) *Weekly energy consumption and variables costs:* We simulate a week of execution of the cloud and determine the corresponding costs. These results only include the variable share of these costs upon which our optimization approach has an influence. This variable share depends on (1) the cost of energy purchased from the additional energy supplier, (2) the PV electricity excess injected into the grid and consumed by other customers, and (3) the network use cost corresponding to the energy exchange between data centers through the virtual pool. The results are shown in Figure 2 and Table II. It must be noted that all these results show the performance on the entire cloud, i.e. including all the data centers.

It can be observed that, when the cloud is managed using SCORPIUS, it consumes significantly less energy than when round-robin is used. More precisely, this difference is equal to

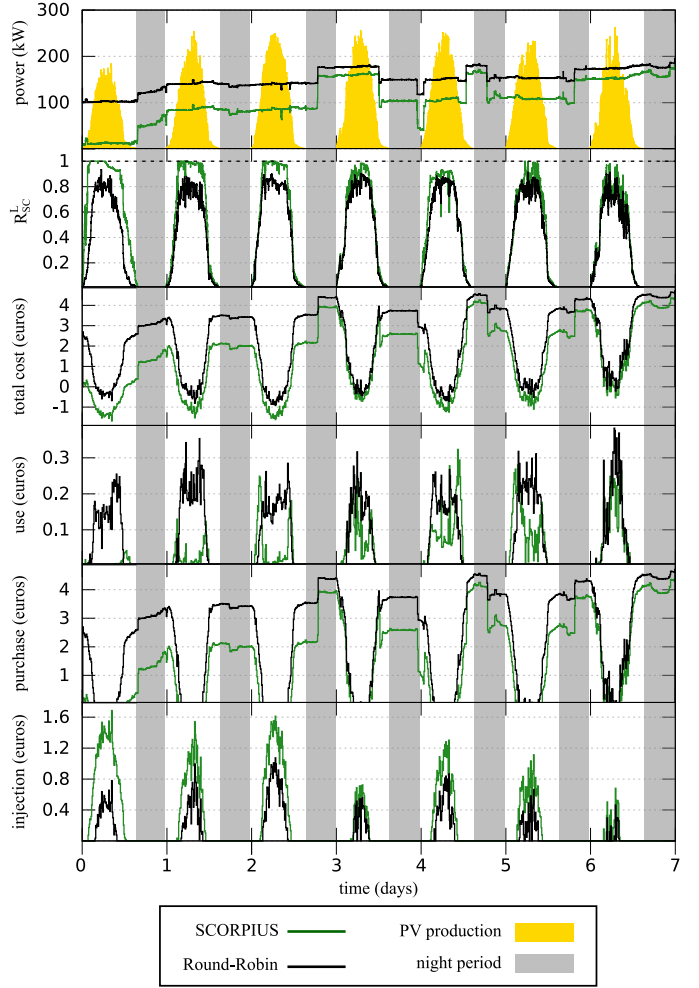


Fig. 2. Cloud performance each time slot over one week.

7.6 MWh over the period of one week considered here. This represents approximately 30% of the cloud weekly energy consumption (relative to the case where round-robin is used). This is the combined effects of switching off unused servers and consolidation through VM migration. Consequently, compared to the case where the round-robin is used, energy purchases from the additional energy supplier are reduced by 5.2 MWh, and energy injections into the grid are increased by 2.3 MWh, leading to a total gain of 959.03 € for one week.

2) *Total annual cost:* The total annual cost include both variable and fixed costs. The variable costs are calculated over the week considered in this study, and they are then multiplied by the number of weeks during a year (52). Monthly fixed costs are multiplied by the number of months in a year (12) to obtain the annual equivalent cost. The total annual costs obtained with the two algorithms considered here are shown in Table II. It must be noted that using SCORPIUS leads to a difference of approximately 50 k€. This amount is extremely significant as it represents around 37% of to the total annual cost with round-robin. Moreover, as SCORPIUS does not require a specific infrastructure, its deployment does

TABLE II
OVERALL CUMULATIVE CLOUD PERFORMANCE.

	total cost		use (week)	purchase (week)	injection (week)	total consumption (week)	self-consumption ratio			
	weekly	annual					days & nights		daytime	
							R_{SC}^L	R_{SC}^C	R_{SC}^L	R_{SC}^C
SCORPIUS	1,446.35 €	83,585.49 €	32.12 € 0.76 MWh	1,640.54 € 10.93 MWh	226.31 € 3.77 MWh	17.65 MWh	34%	38%	54%	61%
round-robin	2,405.38 €	133,455.22 €	63.55 € 1.52 MWh	2,429.1 € 16.19 MWh	87.27 € 1.45 MWh	25.23 MWh	30%	36%	46%	56%
difference	959.03 €	49,869.73 €	31.43 € 0.75 MWh	788.56 € 5.25 MWh	139.04 € 2.31 MWh	7.58 MWh	4 %	2%	7 %	5%

not imply any extra capital cost compared to round-robin.

3) *Cloud self-consumption ratios*: The self-consumption ratios are calculated over each time slot of 10 minutes. The profile of R_{SC}^L is shown on the second row of Figure 2. If both day and night times are taken into account, an average local (resp. cooperative) self-consumption ratio of around 34% (resp. 38%) is obtained when SCORPIUS is used, and of around 30% (resp. 36%) when round-robin is used. However, this self-consumption ratio is equal to zero in both cases during night time when there is no PV electricity generation. Hence, it would be more relevant to consider the performances of both algorithms during day time only. Under these conditions, using SCORPIUS leads to an average local self-consumption ratio of around 54%, while this ratio is only equal to around 46% in the case where round-robin is used. The cooperative self-consumption ratio R_{SC}^C is then equal to around 61% with SCORPIUS and 56% with round-robin.

The local self-consumption ratio increase, in the case when SCORPIUS is used, leads to a reduced need for energy exchange between the data centers. Under the conditions considered here, the amount of energy which is exchanged between the data centers through the virtual pool (and equal to 0.76 MWh) is almost reduced by half with respect to the round-robin case.

V. CONCLUSION AND FUTURE WORK

This paper focuses on the self-consumption optimization of a computing cloud partially powered by PV panels located at each data centers. The study includes economical costs corresponding not only to the cost of energy, but also to the network use cost. This latter cost is based on the network use tariff. An approach for cloud management called SCORPIUS is proposed and compared to the traditionally used and well-known round-robin algorithm. The results show that the increase in terms of self-consumption and total consumption is very significant when the SCORPIUS approach is used. Future work will focus on evaluating our approach with economical models from other countries.

ACKNOWLEDGMENTS

This work has been supported by the Inria exploratory research project COSMIC.

REFERENCES

[1] A. Shehabi, S. Smith, N. Horner, I. Azevedo, R. Brown, J. Koomey, E. Masanet, D. Sartor, M. Herrlin, and W. Lintner, "United States Data Center Energy Usage Report," LBNL, Tech. Rep., 2016.

[2] A. Khosravi, A. Nadjaran Toosi, and R. Buyya, "Online virtual machine migration for renewable energy usage maximization in geographically distributed cloud data centers," *Concurrency and Computation: Practice and Experience*, vol. 29, no. 18, 2017.

[3] T. Dandres, R. F. Moghaddam, K. K. Nguyen, Y. Lemieux, R. Samson, and M. Cheriet, "Consideration of marginal electricity in real-time minimization of distributed data centre emissions," *Journal of Cleaner Production*, vol. 143, pp. 116 – 124, 2017.

[4] European Commission, "Promotion of the use of energy from renewable source, amendments adopted by the european parliament on 17 january 2018 on the proposal for a directive of the european parliament and of the council on the promotion of the use of energy from renewable sources (recast)," European Commission, Tech. Rep., 2018.

[5] Légifrance, "French Republic Official Journal No 0109 of 10th May 2017, text No 22," <https://www.legifrance.gouv.fr/eli/arrete/2017/5/9/DEV1712972A/jo/texte>, 2017.

[6] German Federal Ministry for Economic Affairs and Energy (BMWi), "Mieterstrom: Energiewende im eigenen Haus," <https://www.bmwi.de>, accessed on 18/07/2018.

[7] F. C. Heinrich, T. Cornebize, A. Degomme, A. Legrand, A. Carpen-Amarie, S. Hunold, A. C. Orgerie, and M. Quinson, "Predicting the Energy-Consumption of MPI Applications at Scale Using Only a Single Node," in *IEEE Cluster*, 2017, pp. 92–102.

[8] H. Li, G. Zhu, C. Cui, H. Tang, Y. Dou, and C. He, "Energy-efficient migration and consolidation algorithm of virtual machines in data centers for cloud computing," *Computing*, vol. 98, no. 3, pp. 303–317, Mar 2016.

[9] "Clicking Green: who is winning the race to build a green Internet," Greenpeace report, 2017.

[10] Y. W. Chen and J. M. Chang, "Emaas: Cloud-based energy management service for distributed renewable energy integration," *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2816–2824, Nov 2015.

[11] R. Wolski and J. Brevik, "Using parametric models to represent private cloud workloads," *IEEE Transactions on Services Computing*, vol. 7, no. 4, pp. 714–725, Oct 2014.

[12] "Grid'5000," [urlhttps://www.grid5000.fr](https://www.grid5000.fr).

[13] Y. Li, A.-C. Orgerie, and J.-M. Menaud, "Opportunistic Scheduling in Clouds Partially Powered by Green Energy," in *IEEE Int. Conf. on Green Computing and Communications (GreenCom)*, 2015.

[14] P. Velho, L. Schnorr, H. Casanova, and A. Legrand, "On the Validity of Flow-level TCP Network Models for Grid and Cloud Simulations," *ACM Trans. on Modeling and Comp. Simulation*, vol. 23, no. 4, 2013.

[15] "Photovolta project," <http://photovolta2.univ-nantes.fr>.

[16] M. Boillot, *Advanced Smartgrids for Distribution System Operators*. ISTE and Wiley, 2014.

[17] EDF, "Official website," <https://www.edf.fr/>, Accessed on Jan. 31, 2018.

[18] EDF Energy, "Fixed price electricity contracts," <https://www.edfenergy.com/large-business/buy-energy/fixed-contracts>, Accessed Jan. 31 2018.

[19] OVO Energy, "Official website," <https://www.ovoenergy.com/pay-monthly/2-year-fixed-energy>, Accessed Jan. 31 2018.

[20] "Turpe 5 HBA/BT - Tarifs d'utilisation des réseaux publics de distribution d'électricité," Enedis, http://www.enedis.fr/sites/default/files/plaquette_tarifaire_turpe-5.pdf, 2017.

[21] D. S. Johnson, A. Demers, J. D. Ullman, M. R. Garey, and R. L. Graham, "Worst-case performance bounds for simple one-dimensional packing algorithms," *SIAM J. on Computing*, vol. 3, no. 4, pp. 299–325, 1974.

[22] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, "Versatile, scalable, and accurate simulation of distributed applications and platforms," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2899–2917, Jun. 2014.